

DISTRIBUTED EVENT NOTIFICATION SERVICE

[0001] This invention claims the benefit of US Provisional Application No. 60/273,644, filed March 7, 2001.

5

Field Of The Invention:

[0002] This invention relates to computer environments, and more particularly to event notification in a distributed computing environment.

10 Background and Prior Art

[0003] The Applicant is aware of the following US patents, which serve as prior art in connection with this application:

15 6,195,685, which issued February 27, 2001 to Mukherjee et al;

20 6,185,613, which issued February 6, 2001 to Lawson et al;

25 6,094,681, which issued July 25, 2000 to Shaffer et al; and

5,999,978, which issued December 7, 1999 to Angal et al.

[0004] Event notification and distribution between computer systems connected together within a computer network allows a distributed computing environment to share information and resources. This method of event handling is generally achieved by having applications, which execute on the various computer systems, known herein as event producers (EPs), send events to applications also executing on the various computer systems, known herein as event consumers (ECs). An EP must know which ECs to send to, and ECs must know which EPs to receive from.

[0005] Although current event notification systems provide many benefits, there are several areas that can be improved on. Current event notification systems need the EP to know whether an EC is local (i.e. on the same computer system or 30 subsystem), or whether the EC is remote (i.e. on a different computer system or

subsystem). The EP will carry two separate lists: one for local ECs, and one for remote ECs. Therefore the EP behaves differently depending on whether the EC is local or remote. The same situation and usual implementation is true for the EC; it must know whether the EP is local or remote. An improvement would be to 5 eliminate the requirement for the EP or EC to adapt its behavior based on whether the target EP or EC is local or remote.

[0006] Having local and remote concepts also means that all computer systems must contain accurate, current lists of local and remote EPs and ECs. If an EP or 10 EC were to move its location (for example due to reconfiguration, or a component or network failure), all computer systems must update their local and remote lists before accurate event notification can resume. Automatic coordination of component location would also be an improvement.

[0007] Being able to support EP and/or EC duplication for application redundancy, 15 for example, having the ability to seamlessly transfer responsibilities to another EP or EC in the event of component or network failure, would be an improvement to the art.

[0008] Another deficiency with prior art event notification systems is that, if the 20 event notification is not deliverable (for example due to a network fault that may be temporary in nature), the EP normally receives only a response that the event was undeliverable. A solution that will try repeatedly to deliver an event would be an improvement to the art.

25

Summary of the Invention

[0009] The invention provides a distributed event notification mechanism that allows distributed applications to communicate with each other by generating and receiving the events.

[0010] The invention uses a distributed design rather than a centralized one. Neither a centralized nor a distributed database is required. The invention provides a very simple design for event delivery to the applications that are interested in the events.

5

[0011] The invention adds an EM (Event Manager) function above the EPs and ECs. The purpose of the EM is to abstract EP and EC location information, to provide redundant event operations, and to enhance event notification delivery.

[0012] The invention uses a naming service to find the application that is either the producer or the manager of the event. A naming service can be described, in one particular instance, as a storage database of application names and their locations. The naming service enables network components to connect together without regard for the specific physical locations or configurations of the network.

10
15

[0013] The invention uses a local name and a shared global name for each EM. Each EP uses event lists to maintain a list of ECs that are interested in the event.

20

[0014] The EM can be configured to send an event periodically and repeatedly until an acknowledgment is received, for example, when there are event delivery problems.

25

[0015] The invention allows the EPs and ECs to be moved from one node to another node without impacting the event delivery. It allows the event to be delivered by application name without regard for its location.

[0016] The invention allows for EP and EC applications to have redundant versions to provide event delivery to a standby application when the active application has failed or the network does not permit event delivery.

[0017] Therefore in accordance with a first broad aspect of the invention there is provided a system for event notification and distribution between applications in a distributed computer network comprising: a management function for managing location and event information for distribution between respective applications; and a naming service to enable respective applications to connect together.

5

[0018] In accordance with a second aspect of the invention there is provided a method of providing event notification distribution to applications in a distributed computing system, comprising: providing an event manager to deliver event information to applications having an interest in the information; and using a naming service to locate an application within the system that is either an event producer or an event consumer.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95

Brief Description of the Drawings

[0019] The invention will now be described in greater detail with reference to the attached drawing, wherein Figure 1 is an illustration of an event notification service.

Detailed Description of the Invention

[0020] In a distributed computing environment, many computers/workstations, called nodes, are connected in a network; example, a LAN (Local Area Network). In such a distributed computing environment, different applications may share information and resources. It is necessary for applications to discover events/states occurring in other applications. Each EP maintains a list of ECs that are interested in any event it produces. The EP provides an interface for the EM to access the lists.

20
25

[0021] The invention uses an EM to receive, manage and deliver the events. The EM can be configured to send an event periodically and repeatedly until an

acknowledgment is received. This feature makes the system more robust in the event of a hardware or software failure in the network.

[0022] A distributed naming service that provides name registration and a mechanism for the applications to find a named application is required. Operating systems that provide a naming service are widely available.

[0023] Figure 1 shows the relationships among different applications involved in the event notification process.

[0024] The details of the design are described as follows:

1. Initialization:

[0025] An EM is created at each node, either during node initialization, or as needed following initialization. The EM registers its name with the naming service in the system. The naming service registers a LN (local name) for the local EM, and a GN (Global Name), for all other event managers in the distributed environment. Each EP is also registered with the naming service during initialization, or when it moves from one node to another.

2. Event registration:

[0026] When an EC is interested in an event created by a specific EP, it issues a request to the local EM. The request contains the event name, the EP name and the priority. The EM finds the EP through the naming service. If the EP is local to the node, the EM asks the EP to add the EC to the corresponding event list based on the priority specified in the request. The list is sorted in order of priority. If the EP is not local to the node, the EM sends a request to every EM in the distributed environment by using the GN. The EM on the remote node that contains the

specified EP asks the EP to add the EC to the corresponding event list based on the priority specified in the request.

3. Event Delivery:

5 [0027] When an event is generated, the EP sends the event to the local EM. The EM generates a unique event ID for the event to prevent the same event from being processed twice. The event ID is unique within the entire distributed environment. One approach to creating a unique event ID is to combine a local event sequence number with the local node ID as the event ID. The EM gets the list of the events through an interface provided by the EP, and delivers the event to each EC in the list according to its priority.

4. Moving an EP from one node to another node:

[0028] When an EP is moved from one node to another node (the mechanism to move an application is beyond the scope of this invention), the application takes all event lists with it. It also de-registers itself with the EM in the old node before the move, and re-registers itself with the EM on the new node after the move. When the EP generates an event, it sends the event to the new EM. Event delivery proceeds as described in 3.

20
25 5. Moving an EC from one node to another node:
[0029] Moving an EC (the mechanism to move an application is beyond the scope of this invention) starts with the EC de-registering with the EM in the old node before the move, and re-registering with the EM on the new node after the move.
EC does not re-register the events it has registered with the EM in the old node. If the EC wants to register a new event, it does so with the event manager in the new node.

6. Increasing the robustness of event notification:

[0030] When the EP registers with the EM, it can request that the EM try repeatedly to deliver an event if a delivery problem occurs, before responding that the event was undeliverable. Additionally a limit may be specified for the number of attempted repeats.

5

7. Providing redundant EP:

[0031] The EM can create redundant EPs, possibly to increase the reliability of the network. Because the EC registers only the EPs it is interested in with the local EM, the EM can hide the existence of multiple EPs. Each EP is able to provide the same event service to an EC; service is maintained even in the event of failure of one of the EPs. The mechanism to coordinate the operation of more than one EP is beyond the scope of this invention.

8. Providing redundant EC:

[0032] The EM can create redundant ECs, possibly to increase the reliability of the network. Because the EM provides EC registration and event delivery service on behalf of the EPs, the EP is not concerned with whether there are one or more ECs. The EM can hide the existence of multiple ECs, each EC being able to receive the same event service from an EP. In the event of a failure of one of the ECs, the event information is still successfully received. The mechanism to coordinate the operation of more than one EC is beyond the scope of this invention.

20

[0033] Although particular embodiments of the invention have been described and illustrated it will be apparent to one skilled in the art that numerous changes can be implemented with departing from the basic concept. It is to be understood, however, that such changes will fall within the full scope of the invention as defined by the appended claims.